

THE PROBLEM OF AUTOMATION: INAPPROPRIATE FEEDBACK AND INTERACTION, NOT OVER-AUTOMATION

Donald A. Norman

University of California, San Diego

As automation increasingly takes its place in industry, especially high-risk industry, it is often blamed for causing harm and increasing the chance of human error when failures occur. I propose that the problem is not the presence of automation, but rather its inappropriate design. The problem is that the operations under normal operating conditions are performed appropriately, but there is inadequate feed back and interaction with the humans who must control the overall conduct of the task. When the situations exceed the capabilities of the automatic equipment, then the inadequate feedback leads to difficulties for the human controllers.

The problem, I suggest, is that the automation is at an intermediate level of intelligence, powerful enough to take over control that used to be done by people, but not powerful enough to handle all abnormalities. Moreover, its level of intelligence is insufficient to provide the continual, appropriate feedback that occurs naturally among human operators. This is the source of the current difficulties. To solve this problem, the automation should either be made less intelligent or more so, but the current level is quite inappropriate.

The overall message is that it is possible to reduce error through appropriate design considerations. Appropriate design should assume the existence of error, it should continually provide feedback, it should continually interact with operators in an effective manner, and it should allow for the worst of situations. What is needed is a soft, compliant technology, not a rigid, formal one.

Although automation is often identified as a major culprit in industrial accidents, I propose that the problems result from inappropriate application, not the commonly blamed culprit of 'over-automation'. According to this view, operations would be improved either with a more appropriate form of automation or by removing some existing automation. Current automatic systems have an intermediate level of intelligence that tends to maximize difficulties.

This leads to a second point, namely, that in design, it is essential to examine the entire system: the equipment, the crew, the social structure, learning and training, cooperative activity, and the overall goals of the task. Analyses and remedies that look at isolated segments are apt to lead to local, isolated improvements, but they may also create new problems and difficulties at the system level. Too often, the implementation of some new 'improved' automatic system, warning signal, retraining, or procedure is really a sign of poor overall design: had the proper system level analysis been performed, quite a different solution might have resulted.

I. AUTOMATION: SIMULTANEOUSLY TOO MUCH AND TOO LITTLE

Consider the task of the crew on a modern commercial airplane. Most of the flight activity is routine. Large, modern aircraft are relatively easy to fly: the airplane is stable, responsive, and maneuverable. The automatic equipment

monitors all operations and helps ease the workload of the crew. Indeed, whereas the commercial airplane of a few years ago required a crew of three, the newer planes need only two people to fly them, and most of the time, only one is really necessary. Most of this is good, and the accident rate with modern aircraft has been decreasing over the years, the decrease highly correlated with (and usually thought to be a result of) the introduction of high-technology controls and automation.

There are problems, however. For one, the sheer size of the plane means that the crew cannot know all that is happening. They are physically isolated from the passengers and any difficulties that may be occurring within the passenger section of the plane. They are isolated from most of the physical structures of the aircraft. Even more important than physical isolation is the mental isolation caused by the nature of the controls. The automation tends to isolate the crew from the operations of the aircraft because the automatic equipment monitors and controls the aircraft, providing little or no trace of its operations to the crew, isolating them from the moment-to-moment activities of the aircraft and of the controls. On the one hand, this combination of relative physical and mental isolation from the details of flying helps contribute to the safety by reducing workload and reliance on possible human variability or failure. On the other hand, when the automatic equipment fails, the crew's relative isolation can dramatically increase the difficulties and the magnitude of the problem faced in diagnosing the situation and determining the appropriate course of action.

Physical isolation would be alright if the crew were still up to date on the critical states of the device being controlled. The problem is that, increasingly, the physical isolation is accompanied by a form of mental isolation. Zuboff (1989) describes the control room of a modern paper mill: where once the operators roamed the floor, smelling, hearing and feeling the processes, now they are poised above the floor, isolated in a sound-isolated, air-conditioned, glass control room.

The paper-mill operators do not get the same information about the state of the mill from their meters and displays as they did before from the physical presence. The ship captain does not have a good feel for the actual problems taking place on the other side of the ship. And the automatic equipment in an airplane cockpit can isolate the crew from the state of the aircraft. It is this mental isolation that is thought to be largely responsible for many of the current difficulties.

2. DETECTING SYSTEM PROBLEMS: THREE CASE STUDIES

Here are three case studies from the world of aviation, the domain chosen because aviation is the best documented and validated of all industrial situations.

a. The case of the loss of engine power

In 1985, a China Airlines 747 suffered a slow loss of power from its outer right engine. This would have caused the plane to yaw to the right, but the autopilot compensated, until it finally reached the limit of its compensatory abilities and could no longer keep the plane stable. At that point, the crew did not have enough time to determine the cause of the problem and to take action: the plane rolled and went into a vertical dive of 31,500 feet before it could be recovered. The aircraft was severely damaged and recovery was much in doubt (NTSB 1986; Wiener 1988).

b. The case of the 'incapacitated' pilot

The second case study is a demonstration that lack of information and interaction can take place even in the absence of automation, an important piece of evidence for my argument that automation per se is not the key issue.

In 1979, a commuter aircraft crashed while landing at an airport on Cape Cod, Massachusetts (USA), killing the captain and seriously injuring the first officer and six passengers. The first officer observed that the approach was too low and commented on this to the captain. However, the captain did not respond. But the captain, who was also president of the airline, and who had just hired the first officer, hardly ever responded, even though airline regulations require pilots to do so. Moreover, the captain often flew low. There were obvious social pressures on the first officer that would inhibit further action.

What the first officer failed to notice was that the captain was 'incapacitated', possibly even dead from a heart attack. The U.S. National Transportation Safety Board (NTSB) described it this way:

'The first officer testified that he made all the required callouts except the "no contact" call and that the captain did not acknowledge any of his calls. Because the captain rarely acknowledged calls, even calls such as one dot low (about 50 ft below the 3° glide slope) this lack of response probably would not have alerted the first officer to any physiologic incapacitation of the captain. However, the first officer should have been concerned by the aircraft's steep glidepath, the excessive descent rate, and the high airspeed' (NTSB 1980).

Before you think this a strange, isolated instance, consider this. In an attempt to understand this rather peculiar circumstance, the NTSB noted that United Airlines had earlier performed a study of simulated pilot incapacitation:

'In the United simulator study, when the captain feigned subtle incapacitation while flying the aircraft during an approach, 25 percent of the aircraft hit the "ground". The study also showed a significant reluctance of the first officer to take control of the aircraft. It required between 30 sec and 4 min for the other

crewmember to recognize that the captain was incapacitated and to correct the situation' (NTSB 1980).

c. The case of the fuel leak

In the previous two case studies, the crew was unaware of the developing problems. In this third case study, the vigilant second officer noticed one sign of a problem, but failed to detect another. Here is a quotation from the accident report filed with the NASA Aviation Safety Reporting System (Data Report 64441, dated Feb, 1987). (These are voluntary reports, submitted by the people involved.)

'Shortly after level off at 35,000 ft.... the second officer brought to my attention that he was feeding fuel to all 3 engines from the number 2 tank, but was showing a drop in the number 3 tank. I sent the second officer to the cabin to check that side from the window. While he was gone, I noticed that the wheel was cocked to the right and told the first officer who was flying the plane to take the autopilot off and check. When the autopilot was disengaged, the aircraft showed a roll tendency confirming that we actually had an out of balance condition. The second officer returned and said we were losing a large amount of fuel with a swirl pattern of fuel running about mid-wing to the tip, as well as a vapor pattern covering the entire portion of the wing from mid-wing to the fuselage. At this point we were about 2000 lbs. out of balance....'

In this example, the second officer (the flight engineer) provided the valuable feedback that something seemed wrong with the fuel balance. The automatic pilot had quietly and efficiently compensated for the resulting weight imbalance, and had the second officer not noted the fuel discrepancy, the situation would not have been noted until much later, perhaps too late.

Suppose the automatic pilot could have signaled the crew that it was starting to compensate the balance more than was usual, or at the least, more than when the autopilot was first engaged? This would have alerted the crew to a potential problem. Technically, this information was available to the crew, because the autopilot controls the aircraft by physically moving the real instruments and controls, in this situation, by rotating the control wheel to maintain balance. The slow but consistent turning of the wheel could have been noted by any of the three crew members. This is a subtle cue, however, and it was not noted by either the pilot or the co-pilot (the first officer) until after the second officer had reported the fuel unbalance and had left the cockpit.

3. THE PROBLEM IS NOT AUTOMATION, IT IS LACK OF FEEDBACK

Automation is increasingly blamed for problems in high-risk industry. The general theme of the argument is that in the 'good old days', prior to automation, the controllers were actively engaged in the plant operation. They had to monitor everything and control everything. This had problems, in particular high mental workloads and over-reliance on people's abilities to be continually alert, accurate, and knowledgeable. But it had the virtue of keeping the operators continually informed as to the state of the system.

In the language of control theory or servomechanisms, a system has a desired state, a means for adjusting the system toward that desired state, and then a feedback loop in which the actual state of the system is compared with the desired state, so that additional correction can be performed if there is a mismatch. The combination of this control plus feedback is called the control loop, and when a human is operating the equipment manually, the human is an essential element of the control loop—hence the saying, 'the person is in the loop'.

With the advent of automated controls, the human's job changed. The automation took care of the lower level actions and the human operators simply watched over the system, presumably ever-alert for deviations and problems. Now the human operators were managers or supervisors rather than controllers: they were 'out of the loop' (see the papers in Rasmussen & Rouse 1981, Bainbridge 1987, S. Norman & Orlady 1989, or Weiner & Curry 1980).

Automation has clearly improved many aspects of performance. It leads to superior productivity, efficiency, and quality control. In aircraft, fuel efficiency is improved, schedules can be maintained in inclement weather, and overall accident rates have gone down. But automation also leads to difficulties. When problems arise, the crew may not be sufficiently up-to-date with the current state of the system to diagnose them in reasonable time, and the general reliance on automatic systems may have led to a degradation of manual skills. Finally, the highest stress and workload occur at times of trouble. That is, automatic equipment seems to function best when the workload is light and the task routine: when the task requires assistance, when the workload is highest, this is when the automatic equipment is of least assistance—this is the 'irony' of automation (Bainbridge 1987: see also S. Norman & Orlady 1989).

What of the fact that the people are 'out of the loop'? Is this the major culprit? In some of the case studies in this paper, the crew was clearly out of the loop, failing to detect symptoms of trouble early enough to do anything about them. But in one of the studies, the case of the 'incapacitated' pilot, no automation was involved. Instead, there was an uncommunicative captain, plus social pressures that worked against a junior first officer interrupting the activities of a senior captain. In other words, although the human operators are indeed no longer 'in the loop', the culprit is not automation, it is the lack of continual feedback and interaction.

4. TWO THOUGHT EXPERIMENTS

Consider two thought experiments. In the first, imagine a captain of a plane who turns control over to the autopilot, as in the case studies of the loss of engine power and the fuel leak. In the second thought experiment, imagine that the captain turns control over to the first officer, who flies the plane 'by hand'. In both of these situations, as far as the captain is concerned, the control has been automated: by an autopilot in one situation and by the first officer in the other. But in the first situation, if problems occur, the autopilot will compensate and the crew will notice only by chance (as in the case study of the fuel leak). When automatic devices compensate for problems silently and efficiently, the crew is 'out of the loop', so that when failure of the compensatory equipment finally occurs, they are not in any position to respond immediately and appropriately.

In the case of the second thought experiment where the control was turned over to the first officer, we would expect the first officer to be in continual interaction with the captain. Consider how this would have worked in the case studies of the loss of engine power or the fuel leak. In either case, the problem would almost definitely have been detected much earlier in the flight. The first officer would probably have said something like 'I seem to be correcting this thing more and more—I wonder what's happening?' Yes, from the captain's point of view the rest of the crew serves as a type of automaton, but one that observes and remarks upon conditions. By reporting upon observations and possible discrepancies, each crew member keeps the rest informed and alerted—keeping everyone 'in the loop'.

The observations of these thought experiments are buttressed by the situation described in the case study of the fuel leak, where the second officer, routinely scanning the gauges, noted a puzzling discrepancy and commented on it

to the captain. As the captain's report said, 'the second officer brought to my attention that he was feeding fuel to all 3 engines from the number 2 tank, but was showing a drop in the number 3 tank. I sent the second officer to the cabin to check that side from the window'. Here, even though the second officer did not understand the reason for the discrepant fuel gauge reading, the voiced observation prompted the captain to look over the aircraft by sending the second officer to the cabin to examine the wing and for himself to check the cockpit. The cockpit check led the captain to note that the 'wheel was cocked to the right', which then led to the discovery of the weight imbalance caused by a massive fuel leak. At the time the second officer commented on the fuel gauge reading, he did not know what the problem was, but his comment alerted the crew.

Again, this observation makes the point that the culprit is not actually automation, but rather the lack of feedback. The informal chatter that normally accompanies an experienced, socialized crew tends to keep everyone informed of the complete state of the system, allowing for the early detection of anomalies. Hutchins (in press) has shown how this continual verbal interaction in a system with highly social crews serves to keep everyone attentive and informed, helps the continual training of new members of the crew, and serves as natural monitors for error.

5. THE SOLUTION? MORE APPROPRIATE AUTOMATION

The message is that automation, per se, is not the culprit in high-risk situations. Many of the current problems are indeed a result of automation, but only in the sense that the automation is inappropriately designed and applied.

When people perform actions, feedback is essential for the appropriate monitoring of those actions, to allow for the detection and correction of errors,

and to keep alert. This is hardly a novel point: feedback is an essential aspect of all control theory. But adequate feedback to the human operators is absent far more than it is present, whether the system be a computer operating system, an autopilot, or a telephone system. In fact, it is rather amazing how such an essential source of information could be skipped: the need for complete feedback is one of the major points of Norman (1988). Without appropriate feedback, people are indeed out of the loop: they may not know if their requests have been received, if the actions are being performed properly, or if problems are occurring. Feedback is also essential for learning, both of tasks, and also of the way that the system responds to the wide variety of situations it will encounter.

People construct mental models of systems with which they interact. The model is constructed entirely from what I have called 'the system image', the information available to them from the system, the environment, and their instructions (Norman 1986). But this system image depends critically upon the information displays of modern equipment. When we send a command to an automated piece of equipment, the only way we can update our mental models of the system is through the feedback provided us.

In the first case study, the China Airlines situation where the autopilot kept compensating for the loss of engine power, if the autopilot had been intelligent enough, it might have reported the need to keep compensating. In the case study of the weight imbalance caused by a fuel leak, there were two opportunities to note the problem. An intelligent automaton could have reported on the continual increase in compensation necessary to keep the plane level. Or it might have noted that the fuel level of the number three tank was falling, even though fuel was only supposed to be pumped from the number two tank. And in the case of the incapacitated pilot, if the captain and his first officer had been better socialized and had followed normal and proper callout and response

procedures with the two considered as equal members of the operation, the pilot's incapacitation would have been discovered.

a. We do not know enough to mimic natural human interaction

Note that the problems in all three of the case studies were not due to a lack of information, at least not in the technical sense. Autopilots work by physically moving the same controls that the pilots use. In the case studies of the loss of engine power and the fuel leak, the autopilots compensated by turning the control wheels. In theory, the crew could have noted the problem quite early by noting the position of the wheels, just as the second officer did note an abnormality in the fuel gauge readings in the fuel leak case study. Similarly, there was sufficient information in the case of pilot incapacitation. In these cases the problem was that no person or system commented upon the issues, so that nothing brought the potential problem to the attention of the relevant people. The feedback was potentially available, but it was not attended to properly.

(During the writing of this paper, I took part in an informal replication of the fuel leak incident in the NASA-Ames full-vision, full-motion 727 simulator. Once again, the second officer failed to note the discrepant control wheel position, even though in this case he had read the relevant accident report: the normal cockpit activities drew the focus of attention away from the control wheel position. Our analyses afterwards indicated that the wheel position was not a very salient clue in any case. We plan further studies including a careful replication of this situation as well as a formal experimental study of the two 'thought experiments' described in this paper.)

The task of presenting feedback in an appropriate way is not easy to do. Indeed, we do not yet know how to do it. We do have a good example of how not to inform people of possible difficulties: overuse of alarms. One of the problems of modern automation is the unintelligent use of alarms, each individual

instrument having a single threshold condition that it uses to sound a buzzer or flash a message to the operator, warning of problems. The proliferation of these alarms and the general unreliability of these single-threshold events causes much difficulty (see Patterson 1989, Sorkin 1989, and Sorkin, Kantowitz, & Kantowitz 1988). What is needed is continual feedback about the state of the system, in a normal natural way, much in the manner that human participants in a joint problem-solving activity will discuss the issues among themselves. This means designing systems that are informative, yet non-intrusive, so the interactions are done normally and continually, where the amount and form of feedback adapts to the interactive style of the participants and the nature of the problem. We do not yet know how to do this with automatic devices: current attempts tend to irritate as much as they inform, either failing to present enough information or presenting so much that it becomes an irritant: a nagging, 'back-seat driver', second-guessing all actions.

b. A higher order of awareness is needed

To give the appropriate kind of feedback requires a higher level of sophistication in automation than currently exists. Consider what is required for an automatic pilot to note that it is compensating more than normal. The current automatic systems are feedback loops that attempt to maintain a constant system state. To provide self-monitoring capability that would let it recognize that conditions are changing and more and more compensation is being used, would require a kind of higher-level of awareness, a monitoring of its own monitoring abilities.

Now, obviously, it would not be difficult to build automatic systems for the specific cases of monitoring for increased rudder or control-yoke compensation, or for inappropriate fuel loss: any competent computer-scientist could write an appropriate program. But what about the next problem, one that

will involve yet a different system, yet a slightly different anomaly? We do not know how to solve the general condition.

Consider what would be required of a fuel monitoring system to detect that the fuel level of tank x was dropping, but that fuel was only supposed to be fed from tank y. To solve this problem, in the general case, requires an intelligent system, one that understands the implications of the various control settings of the system. There probably has to be a knowledge base of the systems in the aircraft plus an internal representation for the items that would allow the system to reason about the potential cases. This is the sort of thing done today in laboratories of artificial intelligence and cognitive science, but we do not know how to solve this problem, for the general case. Moreover, even if the automatic monitoring equipment were to note the existence of a system trend or discrepancy that could lead to a difficulty later on, how should it be brought to the attention of the operators in a natural, intelligent fashion, much the way that normal cockpit conversation works?

The solutions will require higher levels of automation, some forms of intelligence in the controls, an appreciation for the proper form of human communication that keeps people well informed, on top of the issues, but not annoyed and irritated. Our current level of knowledge is not enough to do these things.

c. The new irony of over-automation

Many ills have been laid at the feet of 'over-automation'. Too much automation takes the human out of the control loop, it deskills them, and it lowers morale. One much remarked-upon irony of automation is that it fails when it is most needed. I agree with all the analyses of the problems, but from these analyses, I reach the opposite conclusion, a different irony: our current problems with automation, problems that tend to be blamed on 'over-automation', are

probably the result of just the opposite problem—the problem is not that the automation is too powerful, the problem is that it is not powerful enough.

d. Why don't current systems provide feedback?

Why do current systems have such poor feedback and interaction? In part, the reason is a lack of sensitivity on the part of the designer, but in part, it is for a perfectly natural reason: the automation itself doesn't need it! That is, if a designer is asked to design an automatic piece of equipment to control some function, the task is completed when the device functions as requested. Providing feedback and monitoring information to the human operators is of secondary importance, primarily because there does not appear to be any need for it.

Feedback is essential because equipment does fail and because unexpected events do arise. In fact, in any complex task or environment, one should always expect unexpected events: what is unexpected is the type of event that will occur. Human operators need to cope with these situations, and this is why the feedback and 'conversation' is required. Were the equipment never to fail, were it capable of handling all possible situations, then the human operator would not be necessary, so the feedback and interaction would similarly not be necessary. Today, in the absence of perfect automation an appropriate design should assume the existence of error, it should continually provide feedback, it should continually interact with operators in an appropriate manner, and it should have a design appropriate for the worst of situations. What is needed is a soft, compliant technology, not a rigid, formal one.

ACKNOWLEDGMENTS

This research was supported by grant NCC 2-591 to Donald Norman and Edwin Hutchins from the Ames Research Center of the National Aeronautics &

Space Agency in the Aviation Safety/Automation Program. Everett Palmer served as technical monitor. Additional support was provided by funds from the Apple Computer Company and the Digital Equipment Corporation to the Affiliates of Cognitive Science at UCSD. Earlier research has been supported by a grant from the Nippon Telegraph and Telephone Company and the System Development Foundation. I thank Edwin Hutchins for his interactions and guidance and Jim Hollan, Julie Norman, Hank Strub, and Nick Flor for their comments on an earlier version of this manuscript.

REFERENCES

Bainbridge, L. 1987 Ironies of automation. In *New technology and human error* (ed. J. Rasmussen, K. Duncan, & J. Leplat). New York: Wiley.

Hutchins, E. in press. The technology of team navigation. In *Intellectual teamwork: The social and technological foundation of collaborative work* (ed. J. Galegher, R. Kraut, & C. Eijido). Hillsdale, NJ: Lawrence Erlbaum Associates.

Norman, D. A. 1986 Cognitive engineering. In *User centered system design* (ed. D. A. Norman & S. W. Draper, S. W.). Hillsdale, NJ: Lawrence Erlbaum Associates.

Norman, D. A. 1988 *The psychology of everyday things*. New York: Basic Books.

Norman, S. D., & Orlady, H. W. 1989 (Eds.), *Flight deck automation: Promises and realities*. Moffett Field, California: National Aviation and Space Administration.

NTSB. 1980 (National Transportation Safety Board.) Aircraft accident report—Air New England, Inc. DeHavilland DHC-6300, N383 EX Hyannis, Massachusetts. June 17, 1979. (Report No. NTSB/AAR-80/01. Washington: NTSB.

NTSB. 1986 (National Transportation Safety Board.) Aircraft accident report—China Airlines 747-SP, N4522Y, 300 Nautical Miles Northwest of San Francisco, California, February 19, 1985. (Report No. NTSB/AAR-86/03. Washington: NTSB.

NTSB. 1988 (National Transportation Safety Board.) Aircraft accident report—Northwest Airlines, Inc., McDonnell Douglas DC-9-82, N312RC, Detroit Metropolitan Wayne County Airport, Romulus Michigan, August 16, 1987. Report No: NTSB/AAR-88/05, Govt. Accession No: PB 88-91046. May 10, 1988.

Patterson, R. 1989. Auditory warning sounds in the work environment. Paper prepared for the Discussion Meeting, 'Human Factors in High-Risk Situations', The Royal Society (Great Britain), June 28 & 29, 1989, to be published in the *Philosophical Transactions of the Royal Society of London, B*, 1990.

Rasmussen, J., & Rouse, W. B. 1981 (Eds.) *Human detection and diagnosis of system failures*. New York: Plenum Press.

Sorkin, R. D. 1989 Why are people turning off our alarms? *Human Factors Society Bulletin*, 32, No. 4, 3-4.

Sorkin, R. D., Kantowitz, B. K., & Kantowitz, S. C. 1988 Likelihood alarm displays. *Human Factors*, 30, 445-459.

Weiner, E.L. 1988 Cockpit automation. In *Human factors in aviation* (ed. E.L. Weiner & D.C. Nagel). San Diego: Academic Press.

Weiner, E. L., & Curry, R. E. 1980 Flight-deck automation: Promises and Problems. *Ergonomics*, 23, 995-1011. Also in *Pilot error: The human factor* (ed. R. Hurst & L. R. Hurst, 1982. New York: Jason Aronson.

Zuboff, S. 1988 *In the age of the smart machine: The future of work and power*. New York: Basic Books.